REAL Supply

TOPIC 6:
Automation of health and social care

Authors:
Joe Spearing
Nils Gutacker

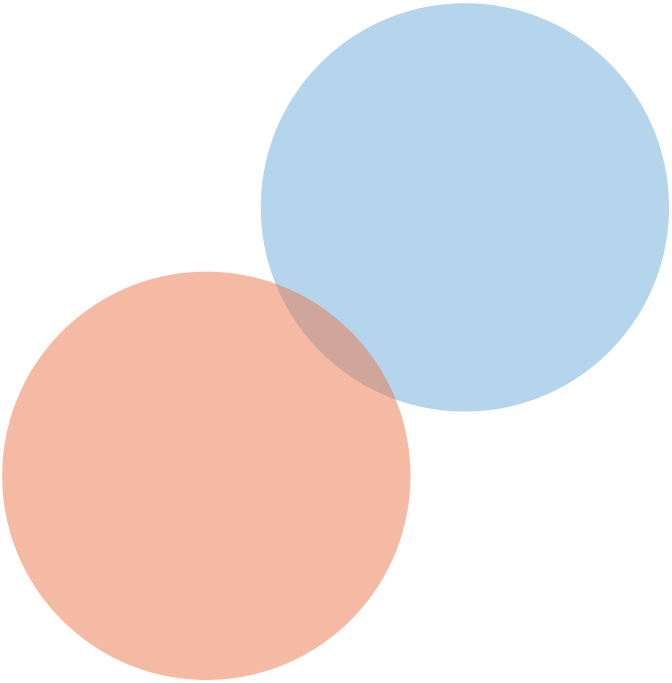

University of York
University of Kent
University of Aberdeen
SPRE
Part of: The Health Foundation REAL Centre

Contents

Plain English summary

Providing health and social care in England is becoming more expensive. A large part of government funding goes toward paying healthcare and social care workers, and even more staff are expected to be hired in the future. But if more money keeps going into these services, there will be less available for other important areas like schools, the police, or the military. So, we need to explore whether some care tasks can be done differently – perhaps in a way that requires fewer workers.

Automation means using technology to do jobs that people usually do. Health and social care have traditionally been hard to automate because they require human skills like decision-making, dexterity, and empathy. However, artificial intelligence (AI) and robots are advancing quickly and changing jobs in many industries. Some experts believe AI and robots could also help control rising costs in health and social care.

In this pathfinder, we examine which health and social care tasks might be automated in the future. We look at past research on automation in other industries and apply those ideas to health and social care. Our goal is to understand what technology can realistically achieve and where human workers remain essential. We also explore how automation might impact employment opportunities and the well-being of workers in this field.

Why study automation?

Public expenditure on health and social care (H&C) in England has grown steadily over the last decades, taking up an increasingly larger share of GDP and government expenditure (Stevenson et al. 2024; Stiebahl 2024). Much of total H&C expenditure is used to pay the wage bill of the workforce, which is driven by two effects: 1) an absolute increase in the size of the health and social care workforce over time; and 2) the so-called 'Baumol effect', by which wages tend to rise faster than labour productivity in labour-intensive industries (Baumol 2012; Baumol and Bowen 1965; Hartwig 2008).[1] These developments threaten the long-term sustainability of the publicly-funded model of H&C provision.[2]

Automation is the process by which capital is substituted for labour to complete a given task. Such automation will, by definition, only occur if the marginal productivity of capital exceeds that of labour so that a unit of care of equivalent quality can be produced cheaper (or more units of equivalent care can be produced at the same cost) when a task is performed by capital rather than labour. Automation is therefore seen as a plausible means to improve the productivity of the health and care system. The probability that automation occurs depends on a range of factors, such as the availability of suitable technologies that permit automation at a sufficiently high level of quality and safety, access to, and availability of, capital and labour inputs in the local factor market, and their associated prices. Economic theory suggests that the scope of automation can be endogenous to rising wages (Cawley, Grabowski, and Hirth 2006; Hémous and Olsen 2022) and demographic pressures (Acemoglu and Restrepo 2021). Finally, it might be determined by institutional features, such as liability regulation and the degree of labour unionisation, managerial performance in allocating capital efficiently, and demand-side responses such as patients' and care recipients' willingness to engage with automated activities.

Recent years have seen increasing public debate about the potential for automation in many sectors of the economy and the downward pressure it may exert on the demand for labour and wages. Such concerns about automation and their impact on employment are not new (e.g. Keynes 1930; Robinson and Acemoglu 2012 and examples cited therein) and have played out in many industries already (Acemoglu and Restrepo 2020). The health and social care sector has traditionally been seen as one where automation is difficult due to the judgement and dexterity involved in largely non-routine tasks. However, the recent emergence of large language models (LLMs)[3] to analyse and detect patterns in big datasets as well as significant improvements in robotics raise the potential that health and social care jobs may become in scope for significant automation over the next decades.

Descriptions of applications of robotics and artificial intelligence (AI), such as LLMs to automate tasks in health and social care as well as assessments of their relative (cost-) effectiveness compared to human labour, are developing at pace (Dvijotham et al. 2023; Hardie et al. 2021; Sandbank et al. 2022; Shin et al. 2023). Indeed, one might argue that the speed of advancement is such that a systematic

1 The productivity of labour-intensive sectors of the economy tends to grow slower than that of more capital-intensive sectors such as manufacturing. Since wages must rise broadly in line with the wider economy, health and social care costs rise in real terms over time, as wages rise faster than productivity growth. This "*cost disease*" is not unique to the English health and care sector, nor to public provision of care in general.

2 For a discussion of the future challenges surrounding recruitment of a sufficient number of health and care workers, see the pathfinder on '**The consequences of non-marginal changes in recruitment into healthcare training**'.

3 ChatGPT was released to the public in November 2022. OpenAI, who develops ChatGPT, was founded in December 2015.

review of existing technologies is quickly outdated and thus not helpful.[4] In this pathfinder, we therefore focus on conceptual issues at a higher level where economic insights may be most useful. This pathfinder explores the potential role of automation in the supply of health and social care in the future, drawing on the wider economic literature on the 'future of work'. We ask which health and care occupations may be at risk of automation. In addition to the technological feasibility of automation, we explore the social, ethical and political limits to automation. Finally, we discuss the potential effects of automation on the health and care workforce, both in terms of demand for their labour and their own well-being.

What is already known?

Economic theories of automation

Economists have made significant progress in understanding the types of roles that are in scope for automation. Much of the recent literature can be traced back to Autor, Levy, and Murnane (2003), who develop a task-based approach to automobility. They define tasks as being cognitive (requiring thought) or manual (requiring physical interaction); and as being routine or non-routine. A task is routine if one could write a set of instructions for doing it. The kinds of tasks which are non-routine require skills such as judgment or creativity. An occupation is then defined by the bundle of tasks it involves.

In Table 1, we provide examples of cognitive, manual, routine and non-routine tasks in health care and non-health care sectors. These examples reveal two things: firstly, there is no necessary relationship between the difficulty involved in performing a task and its routineness. For example, it typically requires more years of education or training before an accountant is trusted to calculate profit from an income statement than before a bricklayer can build a wall. Nevertheless, calculating profit from an income statement is an automatable task because in principle, clear instructions can be written down which obviate the need for judgement or creativity in its execution. Secondly, a key element of non-routine manual tasks is *dexterity*, that is, micro movements designed to adjust to almost imperceptible changes in one's environment. This dexterity is often so intuitive that it cannot be communicated: a nurse cannot say exactly what it is about the "give" of a patient's arm which tells them how hard to drive the needle. This dexterity requirement is seen as a major impediment to automation of certain manual tasks.

4 Those not directly involved in the development and evaluation of new technologies often struggle to appreciate the speed at which these developments occur. An anecdote reported by Frey and Osborne (2017) illustrates this issue well. They write that "*[l]ess than ten years ago [...], Levy and Murnane (2004) pointed at the difficulties of replicating human perception, asserting that driving in traffic is insusceptible to automation [...]. Six years later, in October 2010, Google announced that it had modified several Toyota Priuses to be fully autonomous*" (Frey and Osborne 2017: 255).

Table 1: Examples of cognitive, manual, routine and non-routine tasks in health care and non-health care sectors

	Cognitive	Manual
Routine	Calculate profit from an income statement **Triage appointments**	Printing **Sterilising equipment**
Non-routine	Negotiate a contract **Diagnose an illness**	Lay bricks **Inject a vaccine**

Historically, the range of healthcare tasks in "routine" categories has been limited. While some back-office functions may be automatable, for the most part clinical staff's tasks are generally too complex, either requiring a level of cognitive judgement or dexterity which cannot be translated into a set of rules and communicated to a machine. We will explore a more detailed categorization of roles below.

Automating non-routine tasks

In addition to automating routine tasks, some progress has been made in automating non-routine tasks (Autor 2015). This progress can be put into two categories. Firstly, environmental control standardises variables to avoid the need for dexterous movements. For example, production lines standardise the size and placement of components so that each act of constructing a unit is identical. Applying environmental control in healthcare is fundamentally difficult because, by definition, individual patients differ in ways that cannot be accounted for easily prior to the performance of the task. This feature makes it unlikely that all tasks involved in e.g. general surgery could be automated.

Secondly, advances in artificial intelligence have helped make it possible to automate non-routine cognitive tasks. Supervised machine learning obviates the need to make situational judgments; instead of describing the complex intuition involved in making judgements, a computer algorithm infers optimal decision-making based on observable data. For example, rather than knowing the rules of chess, a machine learning algorithm will analyse data on many past games of chess and identify the "best" move in a statistical sense, i.e., the move that wins in most cases. Alternatively, in principle a machine learning algorithm can infer which patient is most likely to be having a heart attack purely from their observable characteristics and their historical association with heart attacks, rather than any actual medical knowledge. Machine learning therefore has greater potential to help automate non-routine tasks in healthcare.

Estimating automatability using task-based approaches

Given this theoretical framework, a growing economic literature on 'the future of work' (Boyd and Huettinger 2019) has sought to explain which occupations can or will be automated using a 'task-based' approach. Autor, Levy, and Murnane (2003) use this approach to show that the routine task content of US employment has fallen between the 1960s and the late 1990s. However, as we have discussed, using routineness as a predictor of automobility may not accurately predict automation in the future due to advances in AI, which allow for the automation of some non-routine tasks.

To advance beyond this approach, several papers have sought to use the *task content* of occupations to predict which occupations are most susceptible to automation in the future, using different

measures of the kinds of tasks which are automatable. Frey and Osborne (2017) ask experts to classify a small set of occupations according to whether they can be computerized or not (which stands in here for automation in general). They use these categorizations to estimate the probability of computerizability of an occupation as a function of the occupation-specific set of tasks as measured using the US Occupational Information Network (O*NET) data set (provided by the American Bureau of Labour Statistics) and then use the resulting statistical model to predict computerizability for a wider set of occupations. Importantly, the definition of task is at a relatively high level of abstraction; hence information on occupations in some industries are used to estimate the probability of computerizability in other industries. In 2017, Frey and Osborne estimated that 47% of jobs could be automated over the next 20 years.

Arntz, Gregory, and Zierahn (2016) adapt the Frey and Osborne approach by relaxing the assumption that all people who work within the same occupation perform the same tasks. They regress Frey and Osborne's occupation automatability scores on data about the tasks performed at work and then project a person's automation risk from the tasks they actually perform, rather than the average tasks performed in their occupation. They find that the average person within an occupation performs fewer automatable tasks than the median, meaning that occupation level estimates of automatability overestimate the number of people who are exposed to automation. The share of *workers* at risk of automation (defined as 70% of tasks being automated) is therefore lower than the share of *occupations* which are at risk. The Office for National Statistics (2019) apply a similar methodology to the UK, investigating the distribution of automation risk. They find that 25% of health professionals' and 35% of health associate professionals' jobs are at risk of automation.

An alternative to using expert knowledge of which occupations can be automated is to use knowledge of what tasks can be automated. Brynjolfsson and Mitchell (2017) construct a rubric for identifying which tasks can be automated using machine learning, with each task described along 21 properties, each scored from 1 (difficult to automate) to 5 (easy). Brynjolfsson, Mitchell, and Rock (2018) apply this rubric to occupations, finding that many occupations have both a reasonable share of tasks which are susceptible to machine learning and those which are not.

For the most part, these 'future of work' studies have focused on the risk of automation for the full spectrum of occupations in an economy, not health or social care jobs specifically. It is likely that this wide scope comes at a price: the task descriptions used may not be sufficiently granular and specific to provide a good approximation of the tasks performed by health and care workers. Nevertheless, it is instructive to observe how health and social care occupations score in such exercises. Frey and Osborne (2017) include several health care occupations in their data set, a selection of which are presented in Table 2. Notably, there is a large degree of variation in the probability of computerization (and thus automation) in health care occupations in their model.

Table 2: Selected occupations and probability of computerization

Occupation	Probability of Computerization
Mental health and substance abuse social workers	0.0031
Mental health counsellors	0.0048
Health diagnosing and treating practitioners	0.02
Home health aides	0.39
Medical records and health information technicians	0.91

Source: Frey and Osborne (2017)

There is a relevant literature that has focused on administrative processes in health care and how these could be automated. Sahni et al. (2023) analyse several case studies in which machine learning and LLMs can automate aspects of the administration of the US hospital sector. Based on these case studies, they estimate that 5-10% of spending on US health care could be saved through automation and the use of machine learning and LLMs. This work does not speak directly to which occupations may be affected by automation, or any impacts on labour demand more generally (see later sections), but puts a useful boundary on the potential cash savings that may be attainable.

Willis et al. (2020) combine an ethnographic study of primary care workers with the O*NET data set to assess how far primary care work is automatable. The ethnographic survey documents the amount of time spent doing different tasks, which are then categorized and matched to tasks in the O*NET data set. They then estimate how easily automatable each O*NET task is by asking experts to judge the ease of automating a subset of O*NET tasks and then predicting the ease of automating other O*NET tasks using the skills, abilities and knowledge associated with these tasks in the O*NET data set. They conclude that around 44% of administrative tasks within primary care can be automated. Their study is important but limited by it being confined to studying primary care and restricted to only considering which tasks can be automated using existing technology.

How automated is health and social care provision already?

Historically, researchers and practitioners alike have been pessimistic about the potential for health and social care occupations to be automated. In addition to containing a large number of non-routine tasks, many of which are patient-centred and thus reliant on person-to-person interaction, automating regulated roles, especially those involving strong professions, has tended to face potential opposition from powerful interest groups.

To quantitatively explore whether health care occupations are currently more or less automated than non-health care occupations, we turn to the O*NET data set. O*NET provides detailed information about the task content of occupations and their relative importance for this occupation. We use data on tasks, and the degree of automation. The latter variable asks people working within each occupation to assess how automated their occupation already is, on a scale from 1 (‘not at all automated’) to 5 (‘completely automated’). O*NET reports an “average” for each occupation from

between 1 and 5.[5]

Figure 1 plots the distribution of the degree of automation in all occupations, healthcare occupations, and non-healthcare occupations. Notably, the distribution is remarkably similar across healthcare and non-healthcare occupations. The average degree of automation in a healthcare occupation is 2.0, and the average degree of automation in a non-healthcare occupation is 2.1. These quantitative results cast doubt on the idea that automation is less advanced, or less possible, in healthcare than the rest of the economy.

Figure 1: Reported degree of automation in health and non-health care occupations

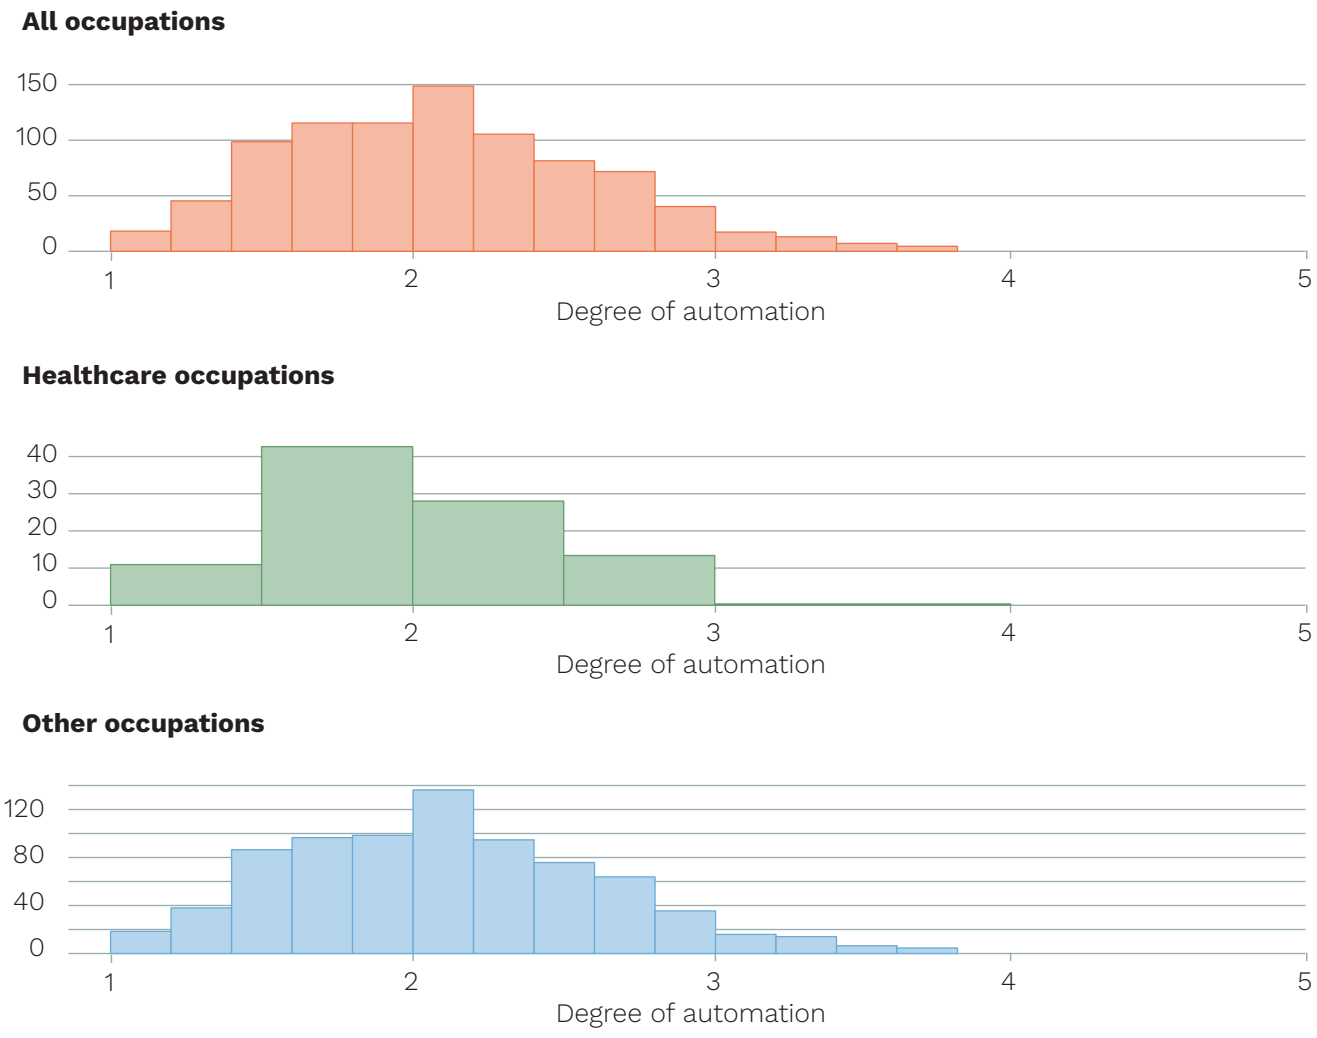

Source: O*NET data set

5 It should be noted that workers in a particular occupation may not be able to accurately assess how far their occupation is automated, particularly as they may not be aware of which automated tasks in their occupation used to be done by humans.

Figure 2: Importance of tasks with differing degree of automatability for selected health care professions

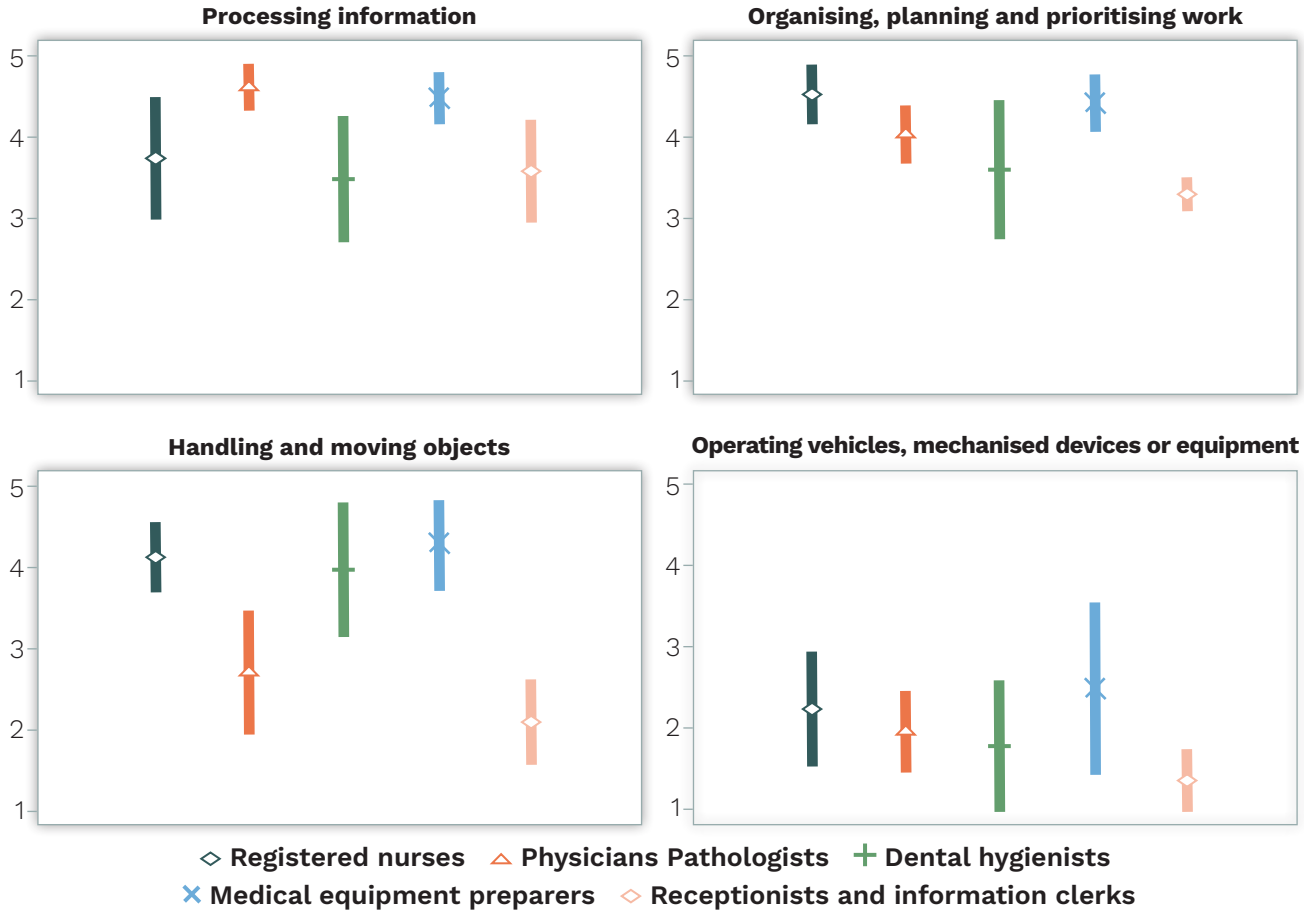

The current task composition of healthcare occupations might have implications for the scope of future potential automation of these occupations. In Figure 2, we plot the importance of a small number of tasks in selected health care occupations. Similar to Autor, Levy, and Murnane (2003), we select tasks which we think are especially susceptible to automation and less susceptible to automation. Processing information might be thought of as relatively easy to automate with modern computer technology, since it often involves relatively routine calculations. LLMs make it easier to quickly extract relevant information from large amounts of data, further increasing the potential to automate data processing. On the other hand, planning and prioritizing requires making value judgments and may include creativity, both aspects of work which may be harder to automate. We also consider two sets of manual tasks. Handling and moving objects requires dexterity, which is a feature of tasks which are hard to automate. On the other hand, operating machinery and equipment might be easier to automate, since it is a potentially repetitive task.

Notably, there are aspects of healthcare occupations which are both easily automatable and not easily automatable. The healthcare occupations we look at score highly for processing information, especially physicians and medical equipment preparers, indicating that part of those occupations might be automatable. Most occupations score relatively low for operating machines, mechanized devices or equipment, however, suggesting that for the most part, these easily automatable manual tasks are not a significant part of healthcare occupations. All our occupations score relatively high for organizing, planning and prioritising work, which indicates a degree of creativity and decision making

which cannot be easily automated. Finally, nurses, hygienists, and medical equipment preparers score highly for dexterous tasks which are unlikely to be easily automated.

Overall, therefore, our analysis of the O*NET data suggests a moderate level of automation has already occurred on average in healthcare occupations, and that in the current task mix of occupations there is more potential for automation.

"A moderate level of automation has already occurred on average in healthcare occupations. But there is more potential for automation."

Limits to automation

Given the potential to automate health and social care tasks, there are several reasons why particular tasks are not automated. An overarching concern is that requisite infrastructure investments will need to be made, which will depend on reform to the NHS capital regime (Darzi 2024).[6] However, in addition, there are a number of considerations which may prevent automation of some tasks. We consider ethical limits, patient preferences, and provider preferences.

There are several ethical constraints on the automation of health and care tasks, especially those that use big data and machine learning. The first relates to safety: policymakers will likely require very high levels of certainty about the efficacy of automated diagnosis, triage or refer tools to be willing to implement them. A significant concern is that the positive results demonstrated using machine learning may not be replicable (Rambachan and Roth 2019). A related concern is inequity created by the use of machine learning and AI: algorithms trained on one population may underperform on another. This concern is particularly acute when it comes to potential racial inequalities. In addition, models trained on real-world data will tend to replicate existing inequalities. Obermeyer et al. (2019) find that, where medical need is estimated using medical costs, models tend to underestimate the medical need of black patients relative to white patients with a similar health status, indicating that supervised learning algorithms interpret unequal access to healthcare as a normative statement about need for healthcare. Until these concerns are addressed, equity will remain a significant impediment to the automation of certain health and care tasks. The second set of ethical concerns relate to the use of data, in particular, norms about the use of patient data for the purposes of training algorithms (Blasimme and Vayena 2020). Regulation of the use of data will need to carefully balance the potential benefits of automation with the risks of misuse of personal data. The final ethical concern is with the dilution of physician responsibility (Spaulding 2020). Delegating diagnosis and medical decision-making to automated systems may be considered unethical because it removes the role of human accountability in the decision-making process.

A second source of resistance to automation may come from patient preferences. Patients are likely to be resistant to any implementation of automation which might be *perceived* as unsafe or unethical even if it is objectively superior. In addition, many healthcare tasks include the maintenance of relationships. Theoretically, they may be part of the "Sweet Talk" economy (McCloskey 2016), where the human relationship is constitutive of the value placed upon the task. For example, patients may place higher value on a nurse reminding them to take their medication than an automated system reminding them

6 Our pathfinder on "**The incentives for long-term investments**" discusses these issues in more detail.

even if they are just as likely to take their medicine. Alternatively, as McCloskey argues in the case of speech therapists, the effectiveness of the intervention might depend on the existence of a human relationship. In particular, some health and social care roles involve empathizing. In these situations, researchers tend to find a strong preference for human relationships as opposed to AI simulations of empathy (Rubin et al. 2024).

"One might also speculate that professions are more likely to 'permit' automation of tasks that they perceive as less valuable to them."

Finally, an interesting feature of healthcare markets is that a large share of the workforce are members of professions, which regulate access to the labour market by enforcing minimum qualification standards and have (restricted) powers to define their own scope of practice. This feature implies that automation that is not in the interests of workers may be less likely to be implemented, because these interest groups will use their power to rule these technologies as unsafe. One might also speculate that professions are more likely to 'permit' automation of tasks that they perceive as less valuable to them (e.g. clinical documentation); either in terms of earning potential or because it limits their ability to carry out other tasks that they find inherently more enjoyable.

Effects of automation on health and care workers

Demand for labour

We understand automation to refer to substituting capital for labour, i.e. having machinery perform a task that was previously performed by a person. Several theoretical and empirical studies have investigated the effect of automation on the demand for labour (e.g. Acemoglu and Autor 2011; Acemoglu et al. 2022; Acemoglu and Restrepo 2018). While these studies focus on the wider labour market, the intuition is likely to apply to health and social care labour markets more specifically as well.

In general, the effect of technology on demand for workers depends on whether technology is a substitute or complement for labour. Automation substitutes capital for labour inputs, lowering the demand for labour in that particular task. It is formally ambiguous whether this increases or decreases demand for workers who perform the tasks which are automated. To see why, note that we can understand the marginal product of a worker as the marginal product of the tasks that worker performs. When tasks are automated there are two offsetting effects on the demand for workers. Firstly, since some tasks no longer require workers to perform them, the workers who are still employed spend a greater percentage of their time performing non-automated tasks. All else, equal, the marginal product of an extra hour spent performing those tasks therefore declines for a given number of workers, reducing demand for those workers. On the other hand, the marginal product of non-automated tasks may depend on which other tasks are performed: if non-automatable and automatable tasks are complements, then an increase in the supply of automatable tasks can raise the marginal product of the non-automatable tasks (Acemoglu and Restrepo 2018; Brynjolfsson, Mitchell, and Rock 2018).

Empirically, there are examples of automation increasing demand for workers and decreasing demand for workers. Bonfiglioli et al. (2024) study the effect of exposure to automation by comparing firms in the same industry with different levels of automatable tasks: exposure to automation is then instrumented using the increase in employment of robots in an industry, multiplied by a firm's share of tasks which can be automated. In their case they find that exposure to automation causes a

decline in employment. Acemoglu and Restrepo (2020) estimate a negative effect of robot adoption on employment on average in the US, whereas Graetz and Michaels (2018) find no negative effect of actual robot use on overall employment. More closely related to this pathfinder is Lee, Iizuka, and Eggleston (2025), who study robot adoption in Japanese nursing homes. They find that robot adoption is associated with an *increase* in labour demand and employment. The association exists only for monitoring robots, i.e. robots that provide real-time data on patients, automating aspects of carers' jobs which can be time consuming, such as monitoring sleep and mobility. These results are at least consistent with the idea that once these tasks have been automated, carers' productivity on the remaining non-automated tasks can increase.

Finally, economists have studied how the effect of automation on employment varies across the type of worker. For example, automation has been shown to affect lower skilled workers more adversely in some contexts (Graetz and Michaels 2018; Lewis 2011). On the other hand, some have argued that where it is possible to automate high earnings jobs (including doctors), automation has the potential to promote greater earnings equality (Agrawal, Gans, and Goldfarb 2023; Susskind and Susskind 2015).

These observations imply that automation, to the extent that it happens, could reconfigure the health and social care workforce in a range of different ways, which are difficult to predict. We might see an increase or decrease in demand for workers, which depends on the complementary of automatable and non-automatable tasks. It is important to point out that a reduction in the demand for labour may not in itself reduce the number of workers employed in the system. The health care labour market is highly regulated, including through central pay negotiation (see 'Pay and working conditions' pathfinder). Current difficulties to recruit and retain staff and associated vacancies imply excess demand for care compared to what the system can deliver. Automation of some tasks may permit workers to carry out tasks that are demanded but could not be delivered beforehand. While this still constitutes a reduction in labour demand, it might not manifest in unemployment.

Effects on workers' wellbeing

The effect that automation of some tasks has on workers' wellbeing is important for considering the likely impact of automation within health and social care. For example, if automation of some tasks has a highly adverse effect on workers' wellbeing, then, given the discussion above, automation is likely to be associated with a contraction in labour supply, meaning that any productivity effects are at least partially offset by an increase in labour costs.

The effects of automating some tasks are likely to depend on how enjoyable the tasks being automated are: to the extent that the tasks which are automated are intellectually not stimulating, repetitive or physically straining, we might expect the result of this automation to be an improvement in the worker's experience. On the other hand, automating easy tasks, or tasks which the worker derives meaning from, may worsen their experience. The theory of automatable tasks does not provide significant help on this front. Routine tasks are neither necessarily easy nor necessarily boring. In addition, as we have discussed, non-routine tasks are increasingly considered automatable, including diagnosis.

As automation drives the dissolution and reconstitution of occupational categories, healthcare professionals may experience task *compression* or task *expansion*. Automation may drive task compression if easier or more routine tasks are automated, leaving a smaller number of workers performing only the tasks which cannot be automated. Task compression can have an adverse effect on workers if the remaining tasks are especially difficult, or if it severely reduces the variety of

tasks performed at work. As an extreme example, we might imagine that machine learning is able to automate the diagnosis of all but the most complex cases, leaving the physician exclusively handling stressful and cognitively taxing cases. Such a scenario would worsen the documented effects of decision fatigue amongst physicians (Persson et al. 2019). On the other hand, task expansion occurs if occupations emerge which specialise in a wide range of non-automatable tasks. Such occupations provide variety for their workers and may be more attractive.

To our knowledge there is little research on the effects of automation on workers' wellbeing. Soffia et al. (2024) collect primary data on UK workers, and find that exposure to robots, AI software, and wearables at work is associated with worse quality of life, while exposure to IT technology was associated with better quality of life. These results might provide suggestive evidence that, overall, automation worsens workers' experience of work. However, the authors are not able to establish causality: it may be that other factors drive this association. Furthermore, the study related to the whole UK economy, rather than health and social care specifically.

Another argument in the literature is about the potential benefits of automation for physical health. Specifically, Gunadi and Ryu (2021) find that exposure to robots decreases the amount of physical tasks performed by unskilled workers in the US and increases their self-reported health, consistent with robots reducing the extent of physical strain on workers. Liu, Luo, and Seamans (2024) find a similar trade-off in the effect of exposure to robots on physical versus mental health in China. Exposure to robots may lead to unhealthy behaviours such as problem drinking (Lu and Fan 2024).

Applied to health and social care in England, we might expect automation to increase stress for workers who primarily perform cognitive tasks, such as doctors. On the other hand, automation may decrease the physical strain on workers such as nurses and care workers who often must perform physical tasks such as lifting and moving patients and clients. Robots which assist lifting, so decrease the number of workers required to lift patients, reduce the amount of time spent performing these physically straining tasks. This instance of automation may improve the health of nurses and enable them to work for longer.

What we do not know

Our survey of the economic literature suggests that while economists have made significant strides in understanding automation, relatively little research has examined its implications for the health and social care sector, particularly in terms of labour demand and worker well-being. Several factors may explain this gap. Historically, as discussed, health and social care have been perceived as sectors with low automation potential. However, as we show in this pathfinder, on some measures health care occupations appear as automated as the average occupation. Additionally, there is a lack of detailed data on the task composition of health and social care roles and the extent to which automation technologies are currently deployed. Much of the existing literature remains focused on individual use cases rather than considering broader system-wide effects on productivity, sustainability, and resilience.

"While economists have made significant strides in understanding automation, relatively little research has examined its implications for the health and social care sector, particularly in terms of labour demand and worker well-being."

This evidence gap leaves several key questions unanswered about the future of health and social care work in England:

1. To what extent does the sector leverage existing automation technologies?

If the NHS and social care system are not yet making optimal use of labour-saving technologies, this could represent a significant opportunity for productivity growth. Addressing this question requires new methodologies to assess both the current level of automation and the potential for further automation in a health and social care context – moving beyond isolated applications to a more comprehensive analysis. Research in this area could help identify which tasks are most amenable to automation, their relative significance within the sector, and the share of employment they represent.

2. What explains any underutilisation of automation technologies?

Our pathfinder on the '**Diffusion of Technology'** explores general barriers to technology adoption, but automation may present unique challenges distinct from other innovations such as pharmaceuticals or surgical techniques. The capital-labour substitution inherent in automation raises concerns about job security, which may, in turn, influence adoption decisions amongst a highly regulated and well-organised labour force. Understanding these dynamics requires further quantitative research to explore institutional, regulatory, and behavioural constraints on diffusion.

3. How quickly do expectations of automation influence training and career choices?

It is unclear how and how fast signals about automation potential feed into decisions around medical training pathways. For instance, has the rise of advanced pattern-recognition technologies in pathology and radiology already influenced the demand for specialty training places? Which medical specialties are most likely to benefit from automation-driven task reallocation, and how might this reshape the future supply of GPs and hospital consultants? Such evidence would help inform the broader debate on whether policymakers should intervene to protect workers from investing in human capital that risks rapid depreciation, or whether individuals already sufficiently internalise these risks into their career decisions and vote with their feet.

4. How will automation affect labour demand and worker well-being?

While we have speculated on automation's likely impact on employment and working conditions, empirical research in an English health and social care context remains scarce. Central to this is an understanding of what tasks are complements or substitutes, and thus whether automation may increase or decrease the demand for labour. A clearer understanding of these relationships would provide valuable input for workforce planning decisions, including the optimal level of future recruitment under the NHS Long Term Workforce Plan (NHS England 2023).

By addressing these questions, future research can contribute to a more evidence-based discussion on the role of automation in shaping the health and social care workforce. Given the growing fiscal pressures on the sector, understanding how technology can enhance productivity while safeguarding worker well-being is an increasingly urgent priority.

What stage of development is the pathfinder at?

Significant uncertainty remains regarding the future automation of health and social care in England and its implications for labour demand and worker well-being. This uncertainty is unsurprising given the rapid pace of technological advancement, which continues to make previously non-automatable tasks increasingly viable for automation. Economics provides a robust framework for analysing both the potential for automation and its broader labour market effects beyond isolated use cases. However, much of the existing research remains theoretical or fails to fully capture the unique characteristics of the health and social care sector. We advocate for more empirical economic analysis to shed light on the complex interactions between technology, labour supply, and labour demand in this sector.

This pathfinder is at an early stage of development and will be iteratively refined as the scientific discourse on automation in health and social care evolves. We plan to expand our discussion of automation to examine its role in changing the skill mix required to perform a given service (see '**Skill-mix and production function**' pathfinder. By automating certain highly skilled and specialised tasks, there may be opportunities to reconfigure service delivery, enabling less specialised and skilled workers to perform the remaining tasks. This, in turn, could reduce overall costs and decrease dependence on highly skilled labour.

References

Acemoglu, D., and Autor, D. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. in David Card and Orley Ashenfelter (eds.), *Handbook of Labor Economics* (Elsevier).

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. 2022. Artificial Intelligence and Jobs: Evidence from Online Vacancies, *Journal of Labor Economics*, 40: S293-S340.

Acemoglu, D., and Restrepo, P. 2018. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment, *American Economic Review*, 108: 1488–542.

———. 2020. Robots and Jobs: Evidence from US Labor Markets, *Journal of Political Economy*, 128: 2188-244.

———. 2021. Demographics and Automation, *The Review of Economic Studies*, 89: 1-44.

Agrawal, A., Gans, J.S., and Goldfarb, A. 2023. Do we want less automation?, *Science*, 381: 155-58.

Arntz, M., Gregory, T., and Zierahn, U. 2016. The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. *OECD Social, Employment and Migration Working Papers*. Paris: OECD Publishing.

Autor, D.H. 2015. Why Are There Still So Many Jobs? The History and Future of Workplace Automation, *Journal of Economic Perspectives*, 29: 3-30.

Autor, D.H., Levy, F., and Murnane, R.J. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration*, *The Quarterly Journal of Economics*, 118: 1279-333.

Baumol, W.J. 2012. *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't* (Yale University Press).

Baumol, W.J., and Bowen, W.G. 1965. On the Performing Arts: The Anatomy of Their Economic Problems, *American Economic Review*, 55: 495-502.

Blasimme, A., and Vayena, E. 2020. 703The Ethics of AI in Biomedical Research, Patient Care, and Public Health. in Markus D. Dubber, Frank Pasquale and Sunit Das (eds.), *The Oxford Handbook of Ethics of AI* (Oxford University Press).

Bonfiglioli, A., Crinò, R., Fadinger, H., and Gancia, G. 2024. Robot Imports and Firm-Level Outcomes, *The Economic Journal*, 134: 3428-44.

Boyd, J.A., and Huettinger, M. 2019. Smithian insights on automation and the future of work, *Futures*, 111: 104-15.

Brynjolfsson, E., and Mitchell, T. 2017. What can machine learning do? Workforce implications, *Science*, 358: 1530-34.

Brynjolfsson, E., Mitchell, T., and Rock, D. 2018. What Can Machines Learn, and What Does It Mean for Occupations and the Economy?, *AEA Papers and Proceedings*, 108: 43–47.

Cawley, J., Grabowski, D.C., and Hirth, R.A. 2006. Factor substitution in nursing homes, *Journal of Health Economics*, 25: 234-47.

Darzi, A. 2024. Independent Investigation of the National Health Service in England.

Dvijotham, K., Winkens, J., Barsbey, M., Ghaisas, S., Stanforth, R., Pawlowski, N., Strachan, P., Ahmed, Z., Azizi, S., Bachrach, Y., Culp, L., Daswani, M., Freyberg, J., Kelly, C., Kiraly, A., Kohlberger, T., McKinney, S., Mustafa, B., Natarajan, V., Geras, K., Witowski, J., Qin, Z.Z., Creswell, J., Shetty, S., Sieniek, M., Spitz, T., Corrado, G., Kohli, P., Cemgil, T., and Karthikesalingam, A. 2023. Enhancing the reliability and accuracy of AI-enabled diagnosis via complementarity-driven deferral to clinicians, *Nature Medicine*, 29: 1814-20.

Frey, C.B., and Osborne, M.A. 2017. The future of employment: How susceptible are jobs to computerisation?, *Technological Forecasting and Social Change*, 114: 254-80.

Graetz, G., and Michaels, G. 2018. Robots at Work, *The Review of Economics and Statistics*, 100: 753-68.

Gunadi, C., and Ryu, H. 2021. Does the rise of robotic technology make people healthier?, *Health Economics*, 30: 2047-62.

Hardie, T., Horton, T., Willis, M., and Warburton, W. 2021. Switched on: How do we get the best out of automation and AI in health care? London: The Health Foundation.

Hartwig, J. 2008. What drives health care expenditure?--Baumol's model of 'unbalanced growth' revisited, *Journal of Health Economics*, 27: 603-23.

Hémous, D., and Olsen, M. 2022. The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality, *American Economic Journal: Macroeconomics*, 14: 179–223.

Keynes, J.M. 1930. Economic Possibilities for our Grandchildren, *The Nation & Athenaeum*: 36-37.

Lee, Y.S., Iizuka, T., and Eggleston, K. 2025. Robots and labor in nursing homes, *Labour Economics*, 92: 102666.

Lewis, E. 2011. Immigration, Skill Mix, and Capital Skill Complementarity*, *The Quarterly Journal of Economics*, 126: 1029-69.

Liu, Q., Luo, S., and Seamans, R. 2024. Pain or anxiety? The health consequences of rising robot adoption in China, *Economics Letters*, 236: 111582.

Lu, W., and Fan, S. 2024. Drinking in despair: Unintended consequences of automation in China, *Health Economics*, 33: 2088-104.

McCloskey, D.N. 2016. *Bourgeois equality: How ideas, not capital or institutions, enriched the world* (University of Chicago Press: Chicago).

NHS England. 2023. NHS Long Term Workforce Plan.

Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. 2019. Dissecting racial bias in an algorithm used to manage the health of populations, *Science*, 366: 447-53.

Office for National Statistics. 2019. Which occupations are at highest risk of being automated?

Persson, E., Barrafrem, K., Meunier, A., and Tinghög, G. 2019. The effect of decision fatigue on surgeons' clinical decision making, *Health Economics*, 28: 1194-203.

Rambachan, A., and Roth, J. 2019. Bias in, bias out? Evaluating the folk wisdom. *arXiv preprint 1909.08518*.

Robinson, J.A., and Acemoglu, D. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty* (Profile Books: London).

Rubin, M., Li, J.Z., Zimmerman, F., Ong, D.C., Goldenberg, A., and Perry, A. 2024. Comparing the Value of Perceived Human versus AI-Generated Empathy.

Sahni, N., Stein, G., Zemmel, R., and Cutler, D.M. 2023. The Potential Impact of Artificial Intelligence on Healthcare Spending, *National Bureau of Economic Research Working Paper Series*, No. 30857.

Sandbank, J., Bataillon, G., Nudelman, A., Krasnitsky, I., Mikulinsky, R., Bien, L., Thibault, L., Albrecht Shach, A., Sebag, G., Clark, D.P., Laifenfeld, D., Schnitt, S.J., Linhart, C., Vecsler, M., and Vincent-Salomon, A. 2022. Validation and real-world clinical application of an artificial intelligence algorithm for breast cancer detection in biopsies, *npj Breast Cancer*, 8: 129.

Shin, H.J., Han, K., Ryu, L., and Kim, E.-K. 2023. The impact of artificial intelligence on the reading times of radiologists for chest radiographs, *npj Digital Medicine*, 6: 82.

Soffia, M., Leiva-Granados, R., Zhou, X., and Skordis, J. 2024. Does technology use impact UK workers' quality of life? A report on worker wellbeing. Institute for the Future of Work.

Spaulding, N.W. 2020. 375Is Human Judgment Necessary?: Artificial Intelligence, Algorithmic Governance, and the Law. in Markus D. Dubber, Frank Pasquale and Sunit Das (eds.), *The Oxford Handbook of Ethics of AI* (Oxford University Press).

Stevenson, G., Bazeer, N., Sameen, H., and Allen, L. 2024. Social care funding: Three key questions about funding in England. London: The Health Foundation.

Stiebahl, S. 2024. NHS funding and expenditure. House of Commons Library.

Susskind, R., and Susskind, D. 2015. *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford University Press).

Willis, M., Duckworth, P., Coulter, A., Meyer, E.T., and Osborne, M. 2020. Qualitative and quantitative approach to assess the potential for automating administrative tasks in general practice, *BMJ Open*, 10: e032412.